


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

 +SRAM +TLB, +(hierarchical +or +hierarchy), +victim +cache

THE ACM DIGITAL LIBRARY
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

Found 8 of

SRAM TLB hierarchical or hierarchy victim cache expand or expanding or extension or extent

148,786

Sort results by

 relevance

 [Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

 expanded form

 [Search Tips](#)
[Try this search in The ACM Guide](#)
 Open results in a new window

Results 1 - 8 of 8

Relevance scale

1 Clock rate versus IPC: the end of the road for conventional microarchitectures

Vikas Agarwal, M. S. Hrishikesh, Stephen W. Keckler, Doug Burger


May 2000 ACM SIGARCH Computer Architecture News, Proceedings of the 27th annual international symposium on Computer architecture, Volume 28 Issue 2
Full text available: [pdf\(207.54 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The doubling of microprocessor performance every three years has been the result of two factors: more transistors per chip and superlinear scaling of the processor clock with technology generation. Our results show that, due to both diminishing improvements in clock rates and poor wire scaling as semiconductor devices shrink, the achievable performance growth of conventional microarchitectures will slow substantially. In this paper, we describe technology-driven models for wire cap ...

2 Multi-level texture caching for 3D graphics hardware



Michael Cox, Narendra Bhandari, Michael Shantz

April 1998 ACM SIGARCH Computer Architecture News, Proceedings of the 25th annual international symposium on Computer architecture, Volume 26 Issue 3
Full text available: [pdf\(1.62 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Traditional graphics hardware architectures implement what we call the *push architecture* for texture mapping. Local memory is dedicated to the accelerator for fast local retrieval of texture during rasterization, and the application is responsible for managing this memory. The push architecture has a bandwidth advantage, but disadvantages of limited texture capacity, escalation of accelerator memory requirements (and therefore cost), and poor memory utilization. The push architecture also ...

3 Energy-aware design of embedded memories: A survey of technologies, architectures, and optimization techniques



Luca Benini, Alberto Macii, Massimo Poncino

February 2003 ACM Transactions on Embedded Computing Systems (TECS), Volume 2 Issue 1
Full text available: [pdf\(288.44 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Embedded systems are often designed under stringent energy consumption budgets, to limit heat generation and battery size. Since memory systems consume a significant amount of energy to store and to forward data, it is then imperative to balance power consumption and performance in memory system design. Contemporary system design focuses on the trade-

off between performance and energy consumption in processing and storage units, as well as in their interconnections. Although memory design is as ...

Keywords: Embedded systems, embedded memories, integration, memories, nonvolatile, system-on-a-chip, volatile

4 UTLB: a mechanism for address translation on network interfaces



Yuqun Chen, Angelos Bilas, Stefanos N. Damianakis, Cezary Dubnicki, Kai Li

October 1998 **Proceedings of the eighth international conference on Architectural support for programming languages and operating systems**, Volume 33 , 32

Issue 11 , 5

Full text available: [pdf\(1.76 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An important aspect of a high-speed network system is the ability to transfer data directly between the network interface and application buffers. Such a *direct data path* requires the network interface to "know" the virtual-to-physical address translation of a user buffer, *i.e.*, the physical memory location of the buffer. This paper presents an efficient address translation architecture, User-managed TLB (UTLB), which eliminates system calls and device interrupts from the common co ...

5 Formal Aspects and Distributed Systems: Functional abstraction driven design space exploration of heterogeneous programmable architectures



Prabhat Mishra, Nikil Dutt, Alex Nicolau

September 2001 **Proceedings of the 14th international symposium on Systems synthesis**

Full text available: [pdf\(195.64 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Rapid Design Space Exploration (DSE) of a programmable architecture is feasible using an automatic toolkit (compiler, simulator, assembler) generation methodology driven by an Architecture Description Language (ADL). While many contemporary ADLs can effectively capture one class of architecture, they are typically unable to capture a wide spectrum of processor and memory features present in DSP, VLIW, EPIC and Superscalar processors. The main bottleneck has been the lack of an abstraction underl ...

Keywords: ADL, DSP, VLIW, design space exploration, functional abstraction, programmable architecture, superscalar

6 Let caches decay: reducing leakage energy via exploitation of cache generational behavior



Zhigang Hu, Stefanos Kaxiras, Margaret Martonosi

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Full text available: [pdf\(873.03 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Power dissipation is increasingly important in CPUs ranging from those intended for mobile use, all the way up to high-performance processors for highend servers. Although the bulk of the power dissipated is dynamic switching power, leakage power is also beginning to be a concern. Chipmakers expect that in future chip generations, leakage's proportion of total chip power will increase significantly. This article examines methods for reducing leakage power within the cache memories of the CPU. Be ...

Keywords: Cache memories, cache decay, generational behavior, leakage power

7 Cache decay: exploiting generational behavior to reduce cache leakage power



Stefanos Kaxiras, Zhigang Hu, Margaret Martonosi

May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual international symposium on Computer architecture**, Volume 29 Issue 2

Full text available:  [pdf\(1.17 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Power dissipation is increasingly important in CPUs ranging from those intended for mobile use, all the way up to high-performance processors for high-end servers. While the bulk of the power dissipated is dynamic switching power, leakage power is also beginning to be a concern. Chipmakers expect that in future chip generations, leakage's proportion of total chip power will increase significantly.

This paper examines methods for reducing leakage power within the cache memori ...

8 Predictor-directed stream buffers



Timothy Sherwood, Suleyman Sair, Brad Calder

December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

Full text available:  [pdf\(187.89 KB\)](#)

 [ps\(1.12 MB\)](#) 

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

[Publisher Site](#)

Results 1 - 8 of 8

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)